TABLE OF CONTENTS

SECTION **PAGE**

Section 14 Online Storage Controller 14-1
14.1 Introduction 14-1
14.2 Storage System 14-1
14.3 Storage Protocol 14-2
14.4 Network Attached Storage Interface 14-4
14.5 Storage Array Network Interface 14-5
14.6 Converged Network Adapter Interface 14-5
14.7 IP Networking 14-5
14.8 Name Services 14-6
14.9 Security Services 14-7
14.10 Interoperability 14-8
14.11 Class of Service and Quality of Service 14-8
14.12 Virtualization 14-9

LIST OF TABLES

TABLE **PAGE**

Table 14.2-1. Replication Operation Modes 14-2
Table 14.6-1. Physical Interfaces for Data Center Bridging 14-5
Table 14.7-1. IP End-to-End Transport Path Models 14-6
Table 14.11-1. Example Storage and Management Protocols 14-8

SECTION 14
ONLINE STORAGE CONTROLLER

14.1 INTRODUCTION

A Data Storage Controller (DSC) is a specialized multiprotocol computer system with an attached disk array that serves in the role of a disk array controller and end node in Base/Post/Camp/Station (B/P/C/S) networks. The DSC is typically a Military Department (MILDEP) asset connected to the Assured Services Local Area Network (ASLAN), but the DSC is not considered part of the ASLAN.

The DSC features and capabilities listed in this section may be offered as part of a unified capability offering associated with other products on the APL. The definitions for DSC are found in Unified Capabilities (UC) Framework 2013, Appendix C, Definitions, Abbreviations and Acronyms, and References.

14.2 STORAGE SYSTEM

DAT-000010 [**Required: DSC**] The system shall provide a Redundant Array of Independent Disks (RAID) for multiple disk drives. The system shall provide a configuration option to select the specific RAID level to be provisioned in the disk array. The RAID levels available for use shall be subject to the specific vendor implementation. At a minimum, the RAID level shall be dual parity RAID-6 for Serial Advanced Technology Attachment (SATA) drives and RAID-5 for Serial Attached Small Computer Systems Interface (SCSI) and Fibre Channel (FC) drives, although stronger RAID levels are acceptable.

DAT-000020 [**Required: DSC**] The system shall be capable of 99.9 percent availability.

DAT-000030 [**Required: DSC**] The system shall provide a management control function for low-level system monitoring and control functions, interface functions, and remote management. The management control function shall provide an Ethernet physical interface(s) for connection to the owner's (i.e., MILDEP) management network/Local Area Network (LAN) and also provide status. The monitoring shall include an initial system check, system cooling fans, temperatures, power supplies, voltages, and system power state tracking and logging.

DAT-000040 [**Required: DSC**] The system shall provide data storage replication (e.g., mirroring) services [Internet protocol (IP) version 4 (IPv4) and version 6 (IPv6)] between systems that are configured as source and destination replication pairs. The replication operations shall provide capabilities for data backup replication, system replication and migration, and system disaster recovery (DR) services in support of continuity of operations (COOP) planning.

DAT-000050 [**Required: DSC**] When the system interfaces to an Integrated Data Protection (IDP) service and the IDP makes copies of data storage information on to another DSC for periodic data storage backup, DR/COOP, migration, and data archiving operation, the system

replication service shall complete the replication regardless of the host connection protocols used between the application servers and the DSC.

DAT-000060 [**Required: DSC**] The system replication and migration services shall provide capabilities to replicate data storage and configuration information onto another standby DSC system for migrating data storage information.

DAT-000070 [**Required: DSC**] The system DR services shall provide capabilities to replicate data storage and configuration information onto another standby DSC system for DR/COOP.

NOTE: This approach provides the threshold capability. Other replication techniques are permitted to ensure communication optimization.

DAT-000080 [**Optional: DSC**] The system shall provide configurable modes for replication (mirroring) operations between the source DSC and the destination DSC. During replication, both the source and the destination must be in a known good state. The configurable modes shall be Asynchronous or Synchronous and are depicted in Table 14.2-1, Replication Operation Modes.

Table 14.2-1. Replication Operation Modes

REPLICATION MODE	DESCRIPTION
Asynchronous (Async)	Incremental, block-based replication between DSCs that occurs as frequently as once per minute by scheduling or manually entering a command to trigger the replication operations.
Synchronous (Sync)	Real-time replication between DSCs that occurs as data is stored or as it changes.

14.3 STORAGE PROTOCOL

DAT-000090 [**Required: DSC**] The system shall provide a Network File System version 3 (NFSv3) server for file systems data input/output (I/O).

DAT-000100 [**Optional: DSC**] The system shall provide a Network File System version 4 (NFSv4) server for file systems data I/O.

DAT-000110 [**Optional: DSC**] The system shall provide a Network File System version 4.1 (NFSv4.1) server, including support for parallel NFS for file systems data I/O.

DAT-000120 [**Required: DCS**] The system shall provide a Common Internet File System version 1.0 (CIFSv1.0) server for file systems data I/O.

DAT-000130 [**Optional: DCS**] The system shall provide a Common Internet File System version 2.0 (CIFSv2.0) server for file systems data I/O.

DAT-000140 [**Optional: DCS**] The system shall provide Internet Small Computer Systems Interface (iSCSI) server (target) operations for data I/O of Logical Units (LUNs) to clients (initiators).

DAT-000150 [**Optional: DCS**] The system shall provide Fibre Channel Protocol (FCP) server (target) operations for data I/O of FCP LUNs to clients (initiators).

DAT-000160 [**Optional: DCS**] The system shall provide Fibre Channel over Ethernet (FCoE) server (target) operations for data I/O of FCP LUNs to clients (initiators).

DAT-000170 [**Optional: DCS**] The system shall provide a HyperText Transfer Protocol Secure (HTTPS) server for file system data I/O and management access to the storage controller operating system. The session shall be secured with Secure Socket Layer (SSL) or Transport Layer Security (TLS), per Internet Engineering Task Force (IETF) Request for Comment (RFC) 5246, and shall comply with Section 4, Information Assurance, for that protocol.

DAT-000180 [**Required: DCS**] The system shall provide Secure Shell version 2 (SSHv2) or SSL for management access to the storage controller operating system. The SSHv2 or SSL implementation shall comply with Section 4, Information Assurance, for that protocol.

DAT-000190 [**Optional: DCS**] The system shall provide Web-based Distributed Authoring and Versioning (WebDAV), per IETF RFC 4918, in support of Cloud-based virtualized storage infrastructures.

DAT-000200 [**Optional: DCS**] The system shall implement the Representational State Transfer (REST) software architecture for distributed hypermedia systems and Cloud-based virtualized storage infrastructures.

DAT-000210 [**Optional: DCS**] The system shall implement the Storage Networking Industry Association (SNIA) Cloud Data Management Interface (CDMI) standard.

DAT-000220 [**Required: DCS**] The system shall provide Global Name Space (GNS) or single name space functionality. The GNS functionality shall provide the capability to aggregate disparate and remote network-based file systems to provide a consolidated view to reduce complexities of localized file management and administration. The GNS functionality shall provide large (i.e., 14 Petabyte [PB] or greater) working pools of disks, transparent data migration, and it shall serve to reduce the number of storage mount points and shares. Each system shall have a dedicated and unique GNS.

NOTE: A GNS functionality is provided with the assumption that it will only be used in deployments where latency is less than 200 ms.

14.4 NETWORK ATTACHED STORAGE INTERFACE

DAT-000230 [**Required: DSC**] The system shall provide physical interfaces for Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10 GbE) services in conformance with Institute of Electrical and Electronics Engineers (IEEE) 802.3 for Ethernet LAN interfaces.

DAT-000240 [**Required: DSC**] The system shall be able to provision, monitor, and detect faults, and to restore Ethernet services in an automated fashion.

DAT-000250 [**Required: DSC**] The system shall provide physical interfaces for out-of-band management (OOBM) access and services with 10/100 Mbps Ethernet interfaces as a minimum. Services shall include remote access with at least one of the following protocols: SSHv2, SSL, HTTPS, and SNMPv3; and the protocols shall be secured in accordance with Section 4, Information Assurance.

DAT-000260 [**Required: DSC**] When the system uses Ethernet, Fast Ethernet, Gigabit Ethernet (GbE), and 10GbE interfaces, the interfaces shall be autosensing, autodetecting, and autoconfiguring with incoming and corresponding Ethernet link negotiation signals.

DAT-000270 [**Required: DSC**] Ethernet services of the system and the Logical Link Interworking Function (IWF) of the system shall terminate the Media Access Control (MAC) layer of Ethernet as described in Ethernet Standard IEEE 802.3.

DAT-000280 [**Required: DSC**] Ethernet services of the system shall support jumbo frames with a configurable Maximum Transmission Unit (MTU) of 9000 bytes or greater, excluding Ethernet encapsulation. When Ethernet encapsulation is included in the frame size calculation, an additional 22 bytes must be included for the MAC header (14 bytes), the Virtual LAN (VLAN) tag (4 bytes), and the Cyclical Redundancy Check (CRC) Checksum (4 bytes) fields in the Ethernet frame, resulting in a maximum of 9022 bytes or greater. The system shall also support a configurable MTU between 1280 bytes and 1540 bytes to ensure packets can transit type 1 encryptors. The system default MTU shall be 1540 bytes.

DAT-000290 [**Required: DSC**] Ethernet services of the system shall allocate a unique Ethernet MAC address to each Ethernet interface associated with a VLAN, as per IEEE 802.1Q.

DAT-000300 [**Required: DSC**] Ethernet services of the system shall support "Link Aggregation," as per IEEE 802.3ad or IEEE 802.1AX-2008, and use with the Link Aggregation Control Protocol.

DAT-000310 [**Optional: DSC**] Ethernet services of the system shall provide Link Layer Discovery Protocol (LLDP), as per IEEE 802.1AB.

14.5 STORAGE ARRAY NETWORK INTERFACE

DAT-000320 [**Optional: DSC**] The system shall provide Fibre Channel (FC) physical interfaces and FCP interfaces and services as per American National Standards Institute (ANSI) X3.230, X3.297, and X3.303.

14.6 CONVERGED NETWORK ADAPTER INTERFACE

DAT-000330 [**Optional: DSC**] The system shall provide physical interfaces for FCoE services over a 10GbE physical interface in conformance with the ANSI T11 FC-BB-5 standard for FCoE with a Converged Network Adapter (CNA).

DAT-000340 [**Optional: DSC**] The system shall provide physical interfaces for Data Center Bridging [DCB, also known as Converged Enhanced Ethernet (CEE)] features, and functionality, per the standards depicted in Table 14.6-1, Physical Interfaces for Data Center Bridging.

Table 14.6-1. Physical Interfaces for Data Center Bridging

DCB STANDARD	DESCRIPTION
IEEE 802.1Qbb for Priority-Based Flow Control (PFC)	Per-Priority PAUSE adds fields to the standard PAUSE frame that allows a device to inhibit transmission of frames on certain priorities as opposed to inhibiting all frame transmissions.
IEEE 802.1Qaz for Enhanced Transmission Selection (ETS)	Enhanced Transmission Selection provides a means for network administrators to allocate link bandwidth to different priorities on the basis of a percentage of total link bandwidth.
IEEE 802.1Qaz Data Center Bridging Exchange Protocol (DCBX)	DCB Exchange is the mechanism in which peers can exchange capabilities to one another with LLDP.
IEEE 802.1Qau for Congestion Notification	Congestion Notification is a mechanism to transmit congestion information on an end-to-end basis per traffic flow.
LEGEND	
DCB: Data Center Bridging	LLDP: Link Layer Discovery Protocol
DCBX: Data Center Bridging Exchange	PFC: Priority-Based Flow Control
ETS: Enhanced Transmission Selection	

14.7 IP NETWORKING

DAT-000350 [**Required: DSC**] The system shall meet the IPv6 requirements defined in Section 5.2.2, Mapping of RFCs to UC Profile Categories, for a simple server/network appliance.

DAT-000360 [**Required: DSC**] The system shall provide statically provisioned or dynamically adjusted large IP packet receive buffers for replication (mirroring) session traffic received on the Ethernet physical interfaces. The receive buffers may be statically provisioned or the operating system of the system may dynamically self-adjust the packet receive buffer size based on

measurements of the E2E path bandwidth, Maximum Segment Size (MSS), Round Trip Time (RTT), and the percentage of packet loss. The system shall provide a default and minimum IP packet receive buffer size of 2048 KB per replication (mirroring) session. The system shall provide a statically provisioned or dynamically adjusting maximum IP packet receive buffer size of up to 8192 KB per replication (mirroring) session.

These IP packet receive buffer size requirements are conceptually based on either the Satellite or Transoceanic and Terrestrial Fiber Optic Cable E2E IP transport path models as depicted in Table 14.7-1, IP End-to-End Transport Path Models.

Table 14.7-1. IP End-to-End Transport Path Models

PATH MODEL	DESCRIPTION
Transoceanic and Terrestrial Fiber Optic Cable	Where an end-to-end terrestrial OC-3 path with 155 Mbps of bandwidth that has an RTT of approximately 250 µs with packet loss of 0.01 percent or less. These characteristics are typical of a transoceanic and terrestrial fiber optic cable path between a pair of cities, such as London and Tokyo. The 2,048 KB buffer size is suitable for these path characteristics.
Satellite	Where an end-to-end satellite DS1 path with 1.544 Mbps of bandwidth that has an RTT of approximately 600 µs with packet loss of 1.0 percent or greater. These characteristics are typical of a satellite path between two locations within the same VSAT footprint. The 8,192 KB buffer size is suitable for these path characteristics.
LEGEND	
DS1: Digital Signal Level 1	µs: Microsecond OC-3: Optical Carrier 3 RTT: Round Trip Time VSAT: Very Small Aperture Terminal
KB: Kilobyte	
Mbps: Megabits per Second	

DAT-000370 [**Required: DSC**] The system shall provide an optimized congestion control (congestion avoidance) algorithm in Transmission Control Protocol (TCP) for avoidance of traffic loss on communications paths in high-speed networks with high latency or large bandwidth-delay products.

NOTE: Two examples of these algorithms currently implemented in modern operating systems are CUBIC TCP in Linux® 2.6.19 and later, and Compound TCP (CTCP) in various Microsoft® operating system products.

14.8 NAME SERVICES

DAT-000380 [**Required: DSC**] The system shall provide Lightweight Directory Access Protocol (LDAP) directory services per IETF RFC 4510.

DAT-000390 [**Required: DSC**] The system shall provide Kerberos authentication service per IETF RFC 4120.

DAT-000400 [**Required: DSC**] The system shall provide Domain Name System (DNS) client functionality.

DAT-000410 [**Required: DSC**] The system shall provide DNS client-side Load Balancing.

DAT-000420 [**Required: DSC**] The system shall provide Network Information Service (NIS) client directory service functionality.

DAT-000430 [**Required: DSC**] The system shall provide NIS Netgroups client directory service functionality.

DAT-000440 [**Optional: DSC**] The system shall provide Network Basic Input/Output System (NETBIOS) over TCP/IP (NBT) Name Resolution and Windows Internet Name Service (WINS).

DAT-000450 [**Required: DSC**] The system shall provide Internet Storage Name Service (iSNS) client functionality per IETF RFC 4171.

DAT-000460 [**Conditional: DSC**] If the system has a Fiber Channel (FC) interface then the system shall provide FC Name and Zone Service.

14.9 SECURITY SERVICES

DAT-000470 [**Optional: DSC**] The system shall provide IPSec per RFC 4301.

DAT-000480 [**Optional: DSC**] The system shall provide Encapsulating Security Payload (ESP) per RFC 4303.

DAT-000490 [**Optional: DSC**] The system shall provide Internet Key Exchange version 2 (IKEv2) per RFC 4306.

DAT-000500 [**Optional: DSC**] The system shall provide a configurable Packet Filter (Firewall) service to block unauthorized access (for intrusion prevention) while permitting authorized communications. The Packet Filter service shall use a "stateless" design that does not degrade performance and shall filter all packets received based on interface, source IP address, protocol, port, Type of Service (TOS), or Time To Live (TTL). The Packet Filter service shall provide a configuration policy for defining combinations of multiple packet match rules and processing actions.

DAT-000510 [**Required: DSC**] The system shall provide encryption of data at rest at a minimum of AES-256 in accordance with Federal Information Processing Standard (FIPS) 140-2 level 1 or higher to provide the following capabilities:

a. Rapid crypto-shredding (destruction) of data, in accordance with National Institute of Standards and Technology (NIST) 800-88, for tactical systems that operate in harm's way and may fall into enemy hands.

b. Rapid recovery from sensitive data spills, where the wrong data is accidentally written to the wrong place.

DAT-000520 [**Required: DSC**] The system shall comply with all appropriate STIGs to include the Database Security Technical Implementation Guide.

14.10 INTEROPERABILITY

DAT-000530 [**Required: DSC**] The system shall provide an Application Programming Interface (API) to enable interaction with other software and systems. The interactions shall include routines, data structures, object classes, and protocols used to communicate between the consumer and implementer of the API. The API protocol and message format (e.g., Extensible Markup Language [XML]) shall be subject to the specific vendor system operating system implementation.

14.11 CLASS OF SERVICE AND QUALITY OF SERVICE

DAT-000540 [**Optoional: DSC**] The system shall provide Class of Service (CoS) and Quality of Service (QoS) marking on egress traffic at layer 2 per IEEE 802.1p and, Section 7.2.1.3, Class of Service Markings, and Section 7.2.1.4, Virtual LAN Capabilities. Traffic classification and marking must occur before the egress transmission of the Ethernet frame with a rule or policy engine that matches on various storage and management protocol types as offered by the system.

NOTE: Examples of Storage Protocols and Management Protocols are listed in Table 14.11-1, Example Storage and Management Protocols.

Table 14.11-1. Example Storage and Management Protocols

STORAGE PROTOCOLS			
NFSv3	NFSv4	NFSv4.1	CIFSv1.0
CIFSv2.0	iSCSI	FCOE	
MANAGEMENT PROTOCOLS			
SSHv2	HTTP/HTTPS/REST	SFTP	SNMP
FTPS	User-defined protocols (e.g., proprietary system to system mirroring protocols		

The marking is made in Ethernet VLAN tags by setting the priority value to between zero and seven, inclusive for various traffic classes. These are to be used in the ASLAN, non-ASLAN, and extended networks for per-hop CoS and QoS traffic conditioning by the network elements.

DAT-000550 [**Required: DSC**] The system shall provide CoS and QoS marking on egress traffic at layer 3 per Section 6, Network Infrastructure End-to-End Performance. Traffic classification and marking must occur before the egress transmission of the IP packet with a rule or policy engine that matches on various storage and management protocols that occur within the system, such as those listed in Table 14.11-1.

NOTE: The IP packets are marked in the TOS field of the IPv6 packet header with Differentiated Services Code Point (DSCP) values from 0 and 63, inclusive. These are to be used in the ASLAN, non-ASLAN, and extended networks for per-hop CoS and QoS traffic conditioning by the network elements.

14.12 VIRTUALIZATION

DAT-000560 [**Optional: DSC**] The system shall provide virtualized Data Storage Controller (vDSC) functionality and individual protocol server processes. The vDSC shall meet all the requirements of a DSC with minor exceptions that are related to design and technical limitations associated with the complete virtualization of an operating system, which include internal counters for attributes of the physical system, QoS traffic processing, and per vDSC Mobile IP correspondent node binding cache limitations.

NOTE: Within the DSC system, a vendor may integrate a third party component(s) that enables virtualization of heterogeneous file servers and provides a GNS capability.

DAT-000570 [**Optional: DSC**] The vDSC capability within the system shall provide secure, Private Networking Domains (PNDs) for Ethernet, VLANs, and IP that isolate the network domains of system units. The PND shall support the use of duplicate IP addresses and IP subnet address ranges among those of any other configured vDSC in the system. The PND shall provide a dedicated IP Forwarding Information Base (FIB) per vDSC.

DAT-000580 [**Optional: DSC**] The vDSC shall provide an individual Command Line Interface (CLI) context with the full command set of the system, with the scope of the commands limited to the individual vDSC CLI context.

DAT-000590 [**Optional: DSC**] The vDSC shall provide a programmatic API with the full command set of the system with the scope of the API commands limited to the individual vDSC context.

DAT-000600 [**Optional: DSC**] The vDSC capability within the system shall provide an individual GNS unique from the system or shall provide a single name space that provides the capability to aggregate disparate hardware and storage architectures into a single file system. The GNS shall provide the capability to aggregate disparate and remote network-based file systems, providing a consolidated view to reduce complexities of localized file management and administration. The GNS shall provide large working pools of disks and transparent data migration, and shall serve to reduce the number of storage mount points and shares. The single name space shall be spread across multiple physical Network Access Server (NAS) heads all representing the same file system without replication. The single name space shall include the ability to tier data automatically within the same file system.